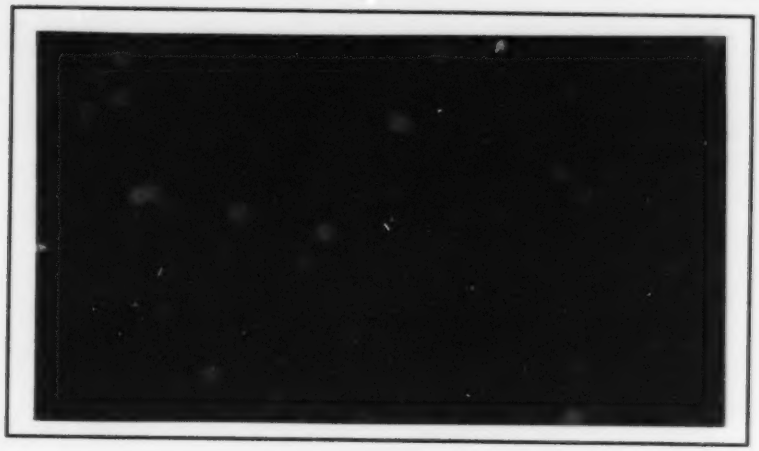

CENTRE FOR
SYSTEMS
SCIENCE

LCCR

Simon Fraser University, Burnaby, B.C., Canada V5A 1S6 / (604) 291-4704

Color Constant Color Indexing

Brian V Funt and Graham D Finlayson

CSS/LCCR TR91-09

October 1991

Color Constant Color Indexing

Brian V. Funt and Graham D. Finlayson
School of Computing Science
Simon Fraser University
Vancouver, B.C.
Canada V5A 1S6
(604) 291-3126
e-mail funt@cs.sfu.ca

Abstract

Objects can be recognized on the basis of their color alone by *color indexing*, a technique developed by Swain and Ballard [14] which involves matching color-space histograms. Color indexing fails, however, when the the incident illumination varies. Although this limitation might be overcome by preprocessing with a color constancy algorithm, we instead propose histogramming color ratios. Since the ratios of color RGB triples from neighbouring locations are relatively insensitive to changes in the incident illumination, this circumvents the need for color constancy preprocessing. Results of tests with the new color-constant-color-indexing algorithm on synthetic and real images show that it works very well.

Keywords: color indexing, color constancy, retinex, object recognition

©1991 B.V. Funt and G.D. Finlayson

92-05827/

1 Introduction

Swain and Ballard [14] developed a very clever, simple scheme that identifies objects entirely on the basis of color. Their method, which they call *color indexing*, radically departs from traditional object recognition strategies based on geometric properties. Color indexing turns out to be remarkably robust in that variations such as a change in orientation, a shift in viewing position, a change in the scene background, partial occlusion or even a radical change in shape (e.g. a shirt tossed onto a chair two different ways) degrade recognition only slightly.

On the other hand, Swain's algorithm is very sensitive to the lighting. Simple changes in the illumination's intensity—let alone its color—radically alter the algorithm's results. Clearly, one solution to this problem is to preprocess the images with some sort of color constancy algorithm [4, 9, 6, 5] in order to remove the effects of varying illumination conditions and, in effect, normalize the images to a standard illuminant. Swain in fact suggests using Novak's [12] supervised color constancy approach in which a few patches of known spectral reflectance are placed in the scene.

While preprocessing for color constancy could well work, it feels unsatisfying in that color constancy appears to be a more difficult problem than color-based object identification. Preprocessing for color constancy destroys the overall elegance inherent in Swain's method.

In this paper, we extend Swain's color indexing method to be illumination independent, not by preprocessing for color constancy, but rather by indexing on an illumination-invariant set of color descriptors. Swain indexes directly on a color triple, we instead index on the Laplacian of the *logarithm* of the colors, which in effect is the ratio of neighbouring colors. Since the illumination remains essentially constant locally, ratioing neighbouring colors removes the illumination component.

Anyone familiar with Land's work will recognize that the roots of this method lie in retinex theory [9]. Our goal here, though, is not color constancy and we do not propose applying full retinex theory. Extending Swain's algorithm requires only illumination-independent descriptors on which to index, it does not require the actual, correct colors. The descriptors we propose, which are in essence the ratios of the correct colors, can be generated much more easily than the correct colors themselves.

We have tested our algorithm on databases of real and synthetic images and compare the results with those of Swain's algorithm. In general, our new algorithm does slightly worse than Swain's under fixed illumination, but substantially better than Swain's (it fails completely) under varying illumination.

2 Summary of Swain's Color Indexing

Intuitively, Swain's color-indexing algorithm identifies an object by comparing its colors to the colors of each object in a database. Crucial to this working, however, is that the total area covered

by each color also be taken into account. The areas are computed and compared by histogramming the images and intersecting the histograms. Color histograms and histogram intersection form the backbone of Swain's method.

A color histogram is three-dimensional and simply represents the count of the number of pixels in an image having a particular color. Before histogramming, Swain converts from camera RGB triples to an opponent-theory-based color space with axes black-white, red-green and blue-yellow. The space is coarsely partitioned along the black-white axis so that histograms which, at 8 bits per channel, potentially could have 256^3 bins have only 2048 bins. Furthermore because the total intensity limits the color differences possible, only 512 of these bins can actually receive counts. Computing color histograms requires time proportional to the number of image pixels, but with special, though not elaborate, hardware Swain reports computation times of 0.04 seconds.

Color histograms of images of every model object, segmented from the image background,are precomputed and stored in a data base. Presented with an image of an unknown object, the color indexing algorithm computes its color histogram and intersects it with every one of the stored histograms in order to find the one that matches best. The intersection of histograms H_a and H_b is defined as:

$$\sum_{i,j,k} max\{H_a(i,j,k), H_b(i,j,k)\} \tag{1}$$

Intersection (or match) values are normalized by the number of pixels in the model histogram, thus matches are between 0 and 1. Histogram intersection only requires time proportional to the number of histogram bins and so is very fast. More sophisticated correlation measures could be used, but Swain's results show that they may not be necessary. Swain reports that the algorithm identifies the correct model most of the time.

3 Illumination Independence

A fundamental assumption of retinex theory and an implication of von Kries adaptation [15] is that the ratio of sensor responses across a color boundary remains unaffected by changes in the incident illumination. This assumption entails a model in which the response of the k^{th} sensor class is formed as the product of scalar coefficients representing illumination and surface albedo relative to that class. For any location x on the surface with incident illumination e_k^x and albedo r_k^x, we have the sensor response ρ_k^x given by

$$\rho_k^x = e_k^x r_k^x \tag{2}$$

This *coefficient model* of sensor response does not hold in general—in fact it is surprising that it holds at all—but it has a long history and for our purposes it is very useful. The main probem with

it is that it depends upon the assumption that for surface reflectances $S_1(\lambda)$ and $S_2(\lambda)$, illumination $E(\lambda)$ and sensor response $R(\lambda)$ the following approximation holds:

$$\frac{\int S_1(\lambda)R(\lambda)d\lambda}{\int S_2(\lambda)R(\lambda)d\lambda} \approx \frac{\int S_1(\lambda)E(\lambda)R(\lambda)d\lambda}{\int S_2(\lambda)E(\lambda)R(\lambda)d\lambda} \tag{3}$$

This is clearly not true in general; however, one way to ensure that the coefficient model holds is to use sensors of narrowband sensitivity. Under the coefficient model, the ratio of sensor responses from two locations under the same illumination yields the ratio of surface albedos:

$$\rho_k^a/\rho_k^b = (r_k^a e_k)/(r_k^b e_k) = r_k^a/r_k^b \tag{4}$$

This ratio has little dependence on illumination other than that it be the same at both locations, which it essentially will, for proximate locations under slowly varying illumination.

Taking logarithms of both sides of (4) turns the ratios into differences

$$\ln(\rho_k^a) - \ln(\rho_k^b) = \ln(r_k^a) - \ln(r_k^b) \tag{5}$$

When these differences are taken between neighbouring pixels in a particular direction, they correspond to differentiation of the logarithm of the image. Hence, under the coefficient model of sensor response, the derivative of the logarithm of the image remains independent of the color and intensity of the incident illumination.

The first directional derivative is non-isotropic, which could lead to orientation affecting object recognition. Natural choices of isotropic operators are the magnitude of the gradient and the Laplacian. We choose the Laplacian, or more precisely the Laplacian of the Gaussian so as to include some smoothing, because it is simpler to compute, it has a theoretical relationship to the center-surround cells of the human visual system [10], and it is in keeping with Horn's [7] lightness algorithm. We do not expect the choice matters much.

Applying the Laplacian to the logarithm of the three channels yields a new 3-tuple for every pixel and it is these that are now histogrammed. Intuitively, in histogramming these Laplacian 3-tuples, we are replacing Swain's measure of color area with a measure of color edge length. In all three channels, the Laplacian will be nearly zero except near the color edges where two regions of different color meet. All pixels along that color edge will generate roughly the same 3-tuple and will accumulate to the same histogram bin. The total accumulation in a particular bin, therefore, represents the length of the boundary between two specific colors.

Even though the color of the light reflected from the two regions changes as the incident illumination changes, their relative color does not. Since the Laplacian tuples reflect the relative color of neighbouring regions, they map to the same histogram bins. Obviously, the illumination does not affect the length of the boundary between regions. The histogram of the Laplacian tuples, therefore, will be invariant to the illumination.

Those Laplacian tuples which are near zero indicate a little bit about the shape/texture of the object. The image of a planar surface under spatially-uniform illumination, for example, would have, in theory, a precisely zero Laplacian except at color edges. A wrinkled surface, on the other hand, would have a wide variety of Laplacian values near zero, but few at zero. Because there is no color change, only a shading change, the Laplacian in all three channels will be identical.

The idea of comparing histograms of Laplacian values specializes quite nicely from three-channel color images to the case of single-band, grey-level images. The texture information continues to be represented in these images and the color edges create albedo changes. The magnitude of the albedo change and the number of pixels at which it occurs remain illumination-independent features that are useful for object recognition. Buchsbaum's [2] analysis shows that the luminance channel contains the majority of the information in color images.

4 The Algorithm

First a database of histograms must be created. The *ratio histogram* of each image is computed as follows:

I. Logarithm step

$$i_k(x,y) \Leftarrow ln(\rho_k(x,y)) \quad k = 1...3$$

II. Laplacian convolution step

$$d_k(x,y) \Leftarrow \nabla^2 G \star i_k(x,y) \quad k = 1...3$$

III. Histogram step

$$H(i,j,k) \Leftarrow \sum_{x,y} \{if\, d_1(x,y) = i \,\&\, d_2(x,y) = j \,\&\, d_3(x,y) = k\, then\, 1\, else\, 0\}$$

Once a database of ratio histograms has been compiled for a set of known objects, it a can be used to recognize an instance of one of these objects in a new image. The ratio histogram of the new image is calculated and then intersected with each of the model histograms in the database. Histogram intersection proceeds according to equation (1), which is exactly as it was for Swain.

Steps I and II represent the only additional computation required to obtain illumination independence. The logarithm in Step I can be done by table lookup in hardware and the Laplacian in Step II is a separable convolution. For this application, the Laplacian requires only a relatively small support since the intention is only to compare colors within small regions so as not to violate the assumption of constant illumination.

4.1 Implementation Details

Since our algorithm operates in logarithm space, the histogram bin distribution should be adjusted accordingly. For ordinary color indexing, histogram bins can be of uniform size with each bin representing an equal volume of color space. A relatively uniform bin distribution follows from an assumed uniform distribution of possible colors. To minimize the effect of variations in illumination intensity, Swain [13] actually samples the black-white axis more coarsely than the other two. Unlike colors, color ratios will not be evenly distributed, which implies that a correspondingly uneven distribution of bin sizes should be used. In particular, the smaller the ratio the more likely its occurrence; thus histogram bins should decrease in size the closer they are to zero, which represents a ratio of one.

Let us consider first the distribution of ratios found in a single color channel. Experimentally, the maximum ratio that occurs is 4.5, so we can assume the ratios fall in the closed interval $[1, 4.5]$ (this could be forced by appropriate scaling). For two intensity values A and B chosen uniformly and randomly from this interval, what is the probability, $\Pr(\max(\frac{A}{B}, \frac{B}{A}) < R)$, that the ratio of $\frac{A}{B}$ (and its reciprocal) is less than R? Rewriting as $\Pr(A < BR) - \Pr(A < \frac{B}{R})$, by standard techniques we can solve for the cumulative probability distribution of the ratios:

$$\Pr(x < R) = \frac{1}{12.25}[21.25 - R - \frac{20.25}{R}] \tag{6}$$

This distribution is graphed in Figure 1, where along the x axis we have marked the bin boundaries which are more frequent nearer the ratio of one. Objects, and therefore images, tend to have large areas of uniform color which causes the histogram bins accumulating ratios near one to dominate during histogram intersection. A ratio of one carries little information about an object's colors, so ratios near one are not included in histogram intersection.

Histogram matches must also be normalized in a slightly different way for ratios than for Swain's color areas. Swain normalizes by the number of pixels in the model histogram; whereas, we normalize to the maximum of the counts in either of the image or model histograms. Since an object is not segmented from the background prior to calculating model histograms, Swain's normalization can result in highly colorful images being falsely matched to a less colorful model. A similar problem would arise if intersections were normalized by the count in image histograms. Normalizing on the maximum of the image and model histograms ensures that a good match occurs only when the intersection is large and the two histograms are of similar size.

While it clearly is more efficient to calculate ratios using logarithms as described above, some of our tests have been conducted by explicitly dividing the values of neighbouring pixels. The logic remains unchanged, however.

5 Test Results

The color constant color indexing algorithm performs well on a variety of real and synthetic images. Objects are correctly identified despite substantial changes in the spectral power distribution of the illuminant. Unsurprisingly, Swain's algorithm performs poorly when the illumination changes. It should be noted that in the tests of Swain's algorithm we use RGB histograms, not opponent-color histograms (he tests both) and prior background segmentation is not performed on the model images.

To evaluate color constant color indexing we first consider whether or not ratios suffice for Swain's original problem under controlled illumination. Second, on synthetic images for which the surface reflectances, illuminants and camera parameters can be completely controlled, we test how the two methods compare. Finally, we test both methods on real images.

5.1 Tests of the Ratio Representation

Even if color ratios are independent of illumination, this says little about ratios as a representation for color indexing. Are ratio histograms sufficient to discriminate between a large number of objects?

To answer this question, we ran the color constant color indexing algorithm on the database of images Swain[1] used in his experiments. First, however, we eliminated 11 of Swain's 66 model images having saturated responses, because ratios relative to saturated pixels cannot be expected to be constant. For our test, then, the model database contains 55 histograms and a second set of 24 different images of the same objects is matched against this database.

Each algorithm's match performance is assessed with reference to three indicators: match rankings, percentile match and match tolerance. The position of the correct match in the sorted list of match values is called its rank, so an image is correctly identified if it has rank 1. The match percentile for each image is defined as $\frac{N-r}{N-1}$, where r is a rank and N is the number of models. Each image is also matched with a certain tolerance relative to the next best matching model. If the correct match has rank i then the match tolerance is $m_i - m_{i-1}$, where m denotes match value. An algorithm that correctly identifies images most of the time, but with high average tolerance, may be preferable to one that correctly identifies images more often, but with lower average tolerance.

Table 1 illustrates the match performance for three algorithms. Swain's, ours with some Gaussian smoothing and ours with no smoothing. Firstly, as Swain reports, color-indexing works well. The second algorithm, color constant indexing with smoothing, shows reasonable performance—19 of the 24 images have 1st place rankings. However, match tolerance is much reduced and, more importantly, two of the images are very poorly matched—ranks of 18 and 27.

The poorer performance can in large part be attributed to the effects of too much smoothing. Swain used reduced images of resolution 128×90, which is quite small relative to the 9×9 Laplacian

[1]The authors are grateful to Michael Swain for providing us with his images.

Algorithm	*No. 1st Rank*	*Other ranks*	*Av. Percentile*	*Av. Tolerance*
Color Indexing	23	2	0.999	0.1212
Laplacian Indexing	19	3,5,18,27	0.961	0.0613
Explicit Ratios	22	2	0.998	0.1023

Table 1: Algorithm Performance : Swain's Images

of Gaussian mask. Under these circumstances, color boundaries will not necessarily be examined in isolation, since the Laplacian operator may straddle more than one edge at at time. To circumvent this problem, we also tried histogramming explicit ratios without smoothing. That is local pairs of pixels are ratioed. Of course, this ratioing can be implemented by a series of directional first-derivative convolutions on the log-color images. The performance for this third algorithm is presented in the third row of Table 1. We conclude that ratios are a sufficiently rich representation for color constant indexing since the ratio histograms perform as well as Swain's color area histograms.

5.2 Tests on Synthetic Images

To the extent that changes in the spectral power distribution of the illumination are modelled by a single scalar multiplication in each sensor channel, the ratio histograms should be relatively illumination independent. To test whether the coefficient rule approximation (Equation 3) holds sufficiently for color ratio indexing, we constructed synthetic images using various measured spectra. These images are free from noise, specularities and other confounding processes that could confuse object identification. As such, they represent a minimal world for object identification.

Surface reflectances were chosen from a set of 40 reflectance spectra that included the Macbeth color checker spectra [11], ceramic spectra [1] and 4 of the natural reflectances (# 98,# 106,# 203,# 324) measured by Krinov [8]. Seven illuminants were used, from reddest to bluest, these are CIE standard A, a 3600K blackbody radiator, D48, D55, D65, D75, D100. Figure 2 shows the sensor response functions derived from the Kodak Wratten filters #25 (red), #58 (green) and #47B (blue) by multiplying by the spectral sensitivity function of our CCD camera. The response of sensor class k, $R_k(\lambda)$, to an illumination $E(\lambda)$ striking a surface reflectance $S(\lambda)$ was calculated as:

$$\rho_k = \int E(\lambda)\, S(\lambda)\, R_k(\lambda)\, d\lambda \tag{6}$$

Thirty synthetic Mondrian objects were generated; a Mondrian contains numerous, two dimensional rectangular patches of uniform reflectance placed side by side. Each Mondrian has the same overall size but contains between 4 and 10 (randomly selected) surface reflectances. If a Mondrian has m patches, then these are distributed according to the formula: patches in x direction = $\lceil \sqrt{m} \rceil$

Algorithm	*No. 1st Rank*	*Failures*	*Av. Percentile*	*Av. Tolerance*
Swain Color Indexing	20	155	N/A	N/A
Color Constant Indexing	180	0	1.000	.568997

Table 2: Algorithm performance : Synthetic images

Images	*No. 1st Rank*	*Other Ranks*	*Failures*	*Av. Percentile*	*Av. Tolerance*
Set1	135	2-5,7,9,10,14,21	10	0.97	0.194
Set2	108	2,3,4,5,9,10	3	0.99	0.256

Table 3: Human Cone Performance

and patches in y direction $= \lceil \frac{m}{\sqrt{m}} \rceil$. Patches are, as far as possible, of uniform size. For example if $m = 7$ then the Mondrian has 3 patches in the first row, 3 in the second, and 1 in the third.

For each illuminant, images of the 30 Mondrians were generated. To separate the issue of brightness change from that of hue change in the illumination, the illuminant spectra were normalized such that their squared area is one. Without loss of generality, the Mondrians imaged under D55 are used as the model set. Match results for Swain's algorithm and for color constant color indexing are given in Table 2. Note the second column displays the number of match failures. An algorithm fails to identify an image if the intersection with the correct model is zero. If this is the case the match rank is undefined.

As expected, Swain's algorithm performs badly—155 of the 180 Mondrians have a zero intersection with the correct model. Indeed, color indexing performs so badly that it is not meaningful to discuss average percentile match or average tolerance. The need for some from of color constancy is readily apparent.

Color constant color indexing performs extremely well. All 180 Mondrians are correctly identified and with high tolerances.

5.2.1 Biological Plausibility

Using the Vos and Walraven[16] estimate of human cone sensitivities as sensors, we can generate synthetic images and examine to what extent color constant color indexing is affected by the choice of cones as sensors. In Table 3 we present the theoretical performance results using the cones. The first row contains the match statistics for all 6 test illuminants—i.e. 180 Mondrians (set1). The second row contains statistics for the test illuminants excluding CIE A and D100 (set2)—120 Mondrians. CIE A and D100 represent the extremes in the spectral variation of the illuminants.

A comparison of Tables 2 and 3 reveals that the broadband nature of the cones does impair the algorithm's performance, but not by too much. Match performance is increased when CIE A and D100—the two extremes of the spectral variation in the illumination—are factored out. Lower

Database	*No. 1st Rank*	*Other Ranks*	*Av. Percentile*	*Av. Tolerance*
3600K	21	2	0.995	0.165
4200K	22		1.000	0.145
5400K	22		1.000	0.137

Table 4: Real Images with Varying Illumination: Color Constant Indexing

Database	*No. 1st Rank*	*Other Ranks*	*Av. Percentile*	*Av. Tolerance*
3600K	14	2,5,7	0.90	0.08
4200K	10	2,3,6,7,8,11	0.768	0.066
5400K	10	2,3,4,8,11	0.80	0.071

Table 5: Real Images with Varying Illumination: Swain's Algorithm

rankings result and both the average match tolerance and average percentile match increase.

5.3 Tests on Real Images

Under three different color temperatures (3600K,4200K and 5400K) pictures were taken of 11 objects comprised of 3 T-shirts, 3 cereal/detergent boxes, 3 sweaters, a Sun User's manual and a child's toy, for a total of 33 images. When the illumination was changed, so were other viewing conditions; shirts and sweaters were deformed, objects were rotated and occluded. The camera responds linearly with intensity and its spectral response functions are as plotted in Figure 2.

Table 4 summarizes the match statistics for color constant color indexing. A model database was constructed using the 11 images taken under one illuminant and then the other 22 images were matched against it. This was repeated for each illuminant. In the table, each row corresponds to a different choice of model database. Performance is good and is independent of the illuminant.

Table 5 tabulates the results for Swain's algorithm. While its performance is poor under varying illumination, it is better than it might have been. This is partly due to the experimental conditions under which the pictures were taken. The color temperature of the illuminant was changed by placing filters in front of the light source. Unfortunately, these filters also diminished the intensity of the light. To compensate for this, camera gain and aperture were adjusted. All pictures were made to have pixels which lie close to the maximum camera response (i.e. 255). Both aperture and gain adjustments are linear so should not affect ratio constancy.

Normalizing images in this way encourages Swain's algorithm to work, since these camera adjustments create an approximate form of color constancy. Nonetheless, even under these favourable experimental conditions Swain's algorithm performs badly. The optimal choice of model set appears to correspond to the 3600K illumination. However, even here 36% of images are wrongly identified; this is extremely poor performance given the small database size. Furthermore, a ranking of 5 or

Database	*No. 1st Rank*	*Other Ranks*	*Av. Percentile*	*Av. Tolerance*
Swain's Images	22	5,8	0.992	0.98
Synthetic Images	180		1.00	0.558
Real Images	19	2	0.986	0.11

Table 6: Matching when Histogram Intersection is a Metric

7 out of 11 is clearly unacceptable.

5.4 Histogram Intersection as a Metric

Swain demonstrates that if two histograms are of the same size then their intersection is a distance metric. In particular, histogram intersection is equivalent to the scaled sum of absolute differences, commonly referred to as the *city-block* metric. Consider the intersection of two histograms M and I each with n bins.

$$if \sum_{i=1}^{n} M_i = \sum_{i=1}^{n} I_i = T$$

$$\Rightarrow 1 - H(I,M) = \frac{1}{2T} \sum_{i=1}^{n} |I_i - M_i|$$

Since ratios close to one are ignored ratio histograms must be normalised to have equal total bin counts. Results for color constant color indexing when this metric condition is enforced are given in Table 6.

It appears that color constant color indexing continues to work well. There is a slight performance fall; however, this may not be surprising considering the type of normalization. The results in Table 6 illustrate the stability of ratio histograms as a context-invariant object descriptor. Since normalized histograms have constant size, this stability suggests match performance could be increased by using a K-nearest neighbour classifier. For a discussion of these classifiers see Duda and Hart[3].

The model database for a K-nearest neighbour classifier contains the ratio histograms of objects imaged in many different visual contexts. An image is identified by examining the K best matches in this duplicate database. Of course, if all K matches are of the same object then this is a strong match. In general, however, it is sufficient to select the most numerous matched object as the identity of an image. A K-nearest neighbour classifier requires that ratio histograms be a stable representation and that their intersection be a metric. We predict, but have yet to test, that the performance of color constant color indexing will be improved if a K nearest neighbour classifier is used. Swain's method cannot be extended in this way, since color histograms are not stable under illumination change.

6 Conclusion

Color constant color indexing based on comparing the color of neighbouring locations works very well. Our experiments show that it extends Swain and Ballard's color-based object recognition scheme to scenes of uncontrolled illumination without affecting overall performance. The strategy of histogramming color ratios (in essence edges) instead of areas circumvents the need for an illumination-independent description of the actual colors in the scene.

7 Acknowledgement

This work has been supported by the Centre for Systems Science at Simon Fraser University and the Natural Sciences and Engineering Research Council of Canada. We also wish to thank Michael Swain for providing us with his image database.

8 Figure captions

Figure 1. The cumulative probability distribution for ratios when all colors are equally likely and the maximum ratio is 4.5.

Figure 2. The camera response function of our CCD camera multiplied by the Kodak Wratten filters numbers #25, #47B and #58.

References

[1] R.S. Berns and K.H. Petersen. Empirical modelling of systematic spectrophotometric errors. *Color Res. Appl.*, (4):243, 1988.

[2] G. Buchsbaum and A. Gottschalk. Trichromacy, opponent colours coding and optimum colour information transmission in the retina. *Proc. R. Soc. Lond. B.*, 220:89–113, 1983.

[3] R.O. Duda and P.E. Hart. *Pattern Classification and Scene Analysis.* John Wiley and Sons, 1973.

[4] D. Forsyth. A novel algorithm for color constancy. *Int. J. Comput. Vision*, 5:5–36, 1990.

[5] B.V. Funt and M.S. Drew. Color constancy computation in near–Mondrian scenes using a finite dimensional linear model. In *Computer Vision and Pattern Recognition Proceedings*, pages 544–549. IEEE Computer Society, June 1988.

[6] B.V. Funt and J. Ho. Color from black and white. In *Proceedings of the Second International Conference on Computer Vision, Tarpon Springs Dec 5-8*, pages 2–8. IEEE, 1988. and *Int. J. Computer Vision*, 3:109-117, 1989.

[7] B. K. P. Horn. Determining lightness from an image. *Computer Vision, Graphics, and Image Processing*, 3:277–299, 1974.

[8] E.L. Krinov. Spectral reflectance properties of natural formations. *Technical Translation TT-439, National Research Council of Canada*, 1947.

[9] E.H. Land and J.J. McCann. Lightness and retinex theory. *J. Opt. Soc. Amer.*, 61:1–11, 1971.

[10] D. Marr. *Vision.* Freeman, 1982.

[11] C.S. McCamy, H. Marcus, and J.G. Davidson. A color-rendition chart. *J. App. Photog. Eng.*, pages 95–99, 1976.

[12] C.L. Novak and S.A. Shafer. Supervised color constancy using a color chart. Technical Report CMU-CS-90-140, Carnegie Mellon University School of Computer Science, 1990.

[13] M.J. Swain. *Color Indexing (PhD Thesis).* University of Rochester Department of Computer Science, 1990.

[14] M.J. Swain and D.H. Ballard. Indexing via color histograms. In *Proceedings: International Conference on Computer Vision, Osaka , Dec.4-7/90*, pages 390–393. IEEE, 1990.

[15] G. West and M.H. Brill. Necessary and sufficient conditions for von Kries chromatic adaption to give colour constancy. *J. Math. Biol.*, 15:249–258, 1982.

[16] G. Wyszecki and W.S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas.* Wiley, New York, 2nd edition, 1982.

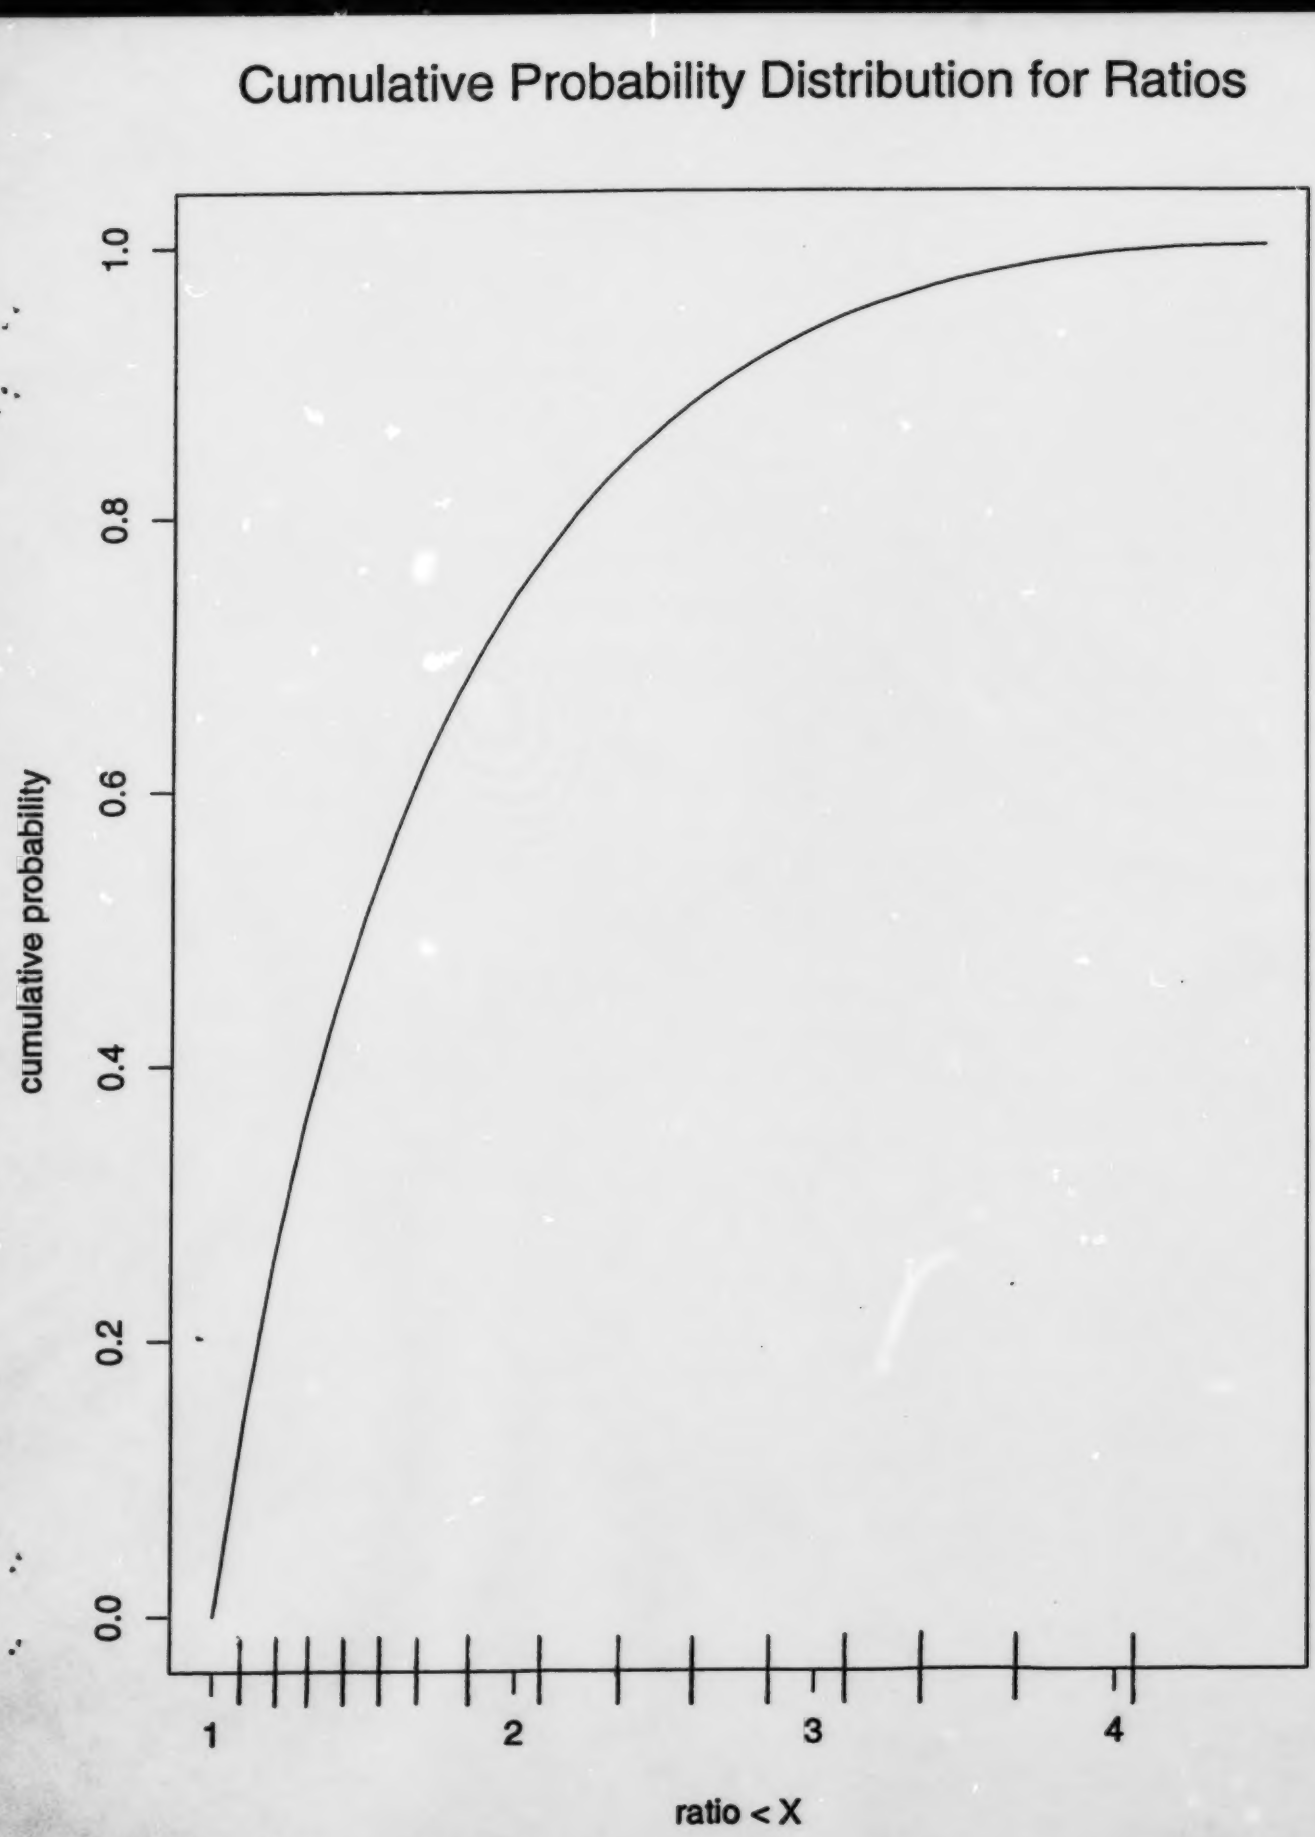
Cumulative Probability Distribution for Ratios
cumulative probability
0.0
0.2
0.4
0.6
0.8
1.0
1
2
3
4
ratio < X

Figure 1

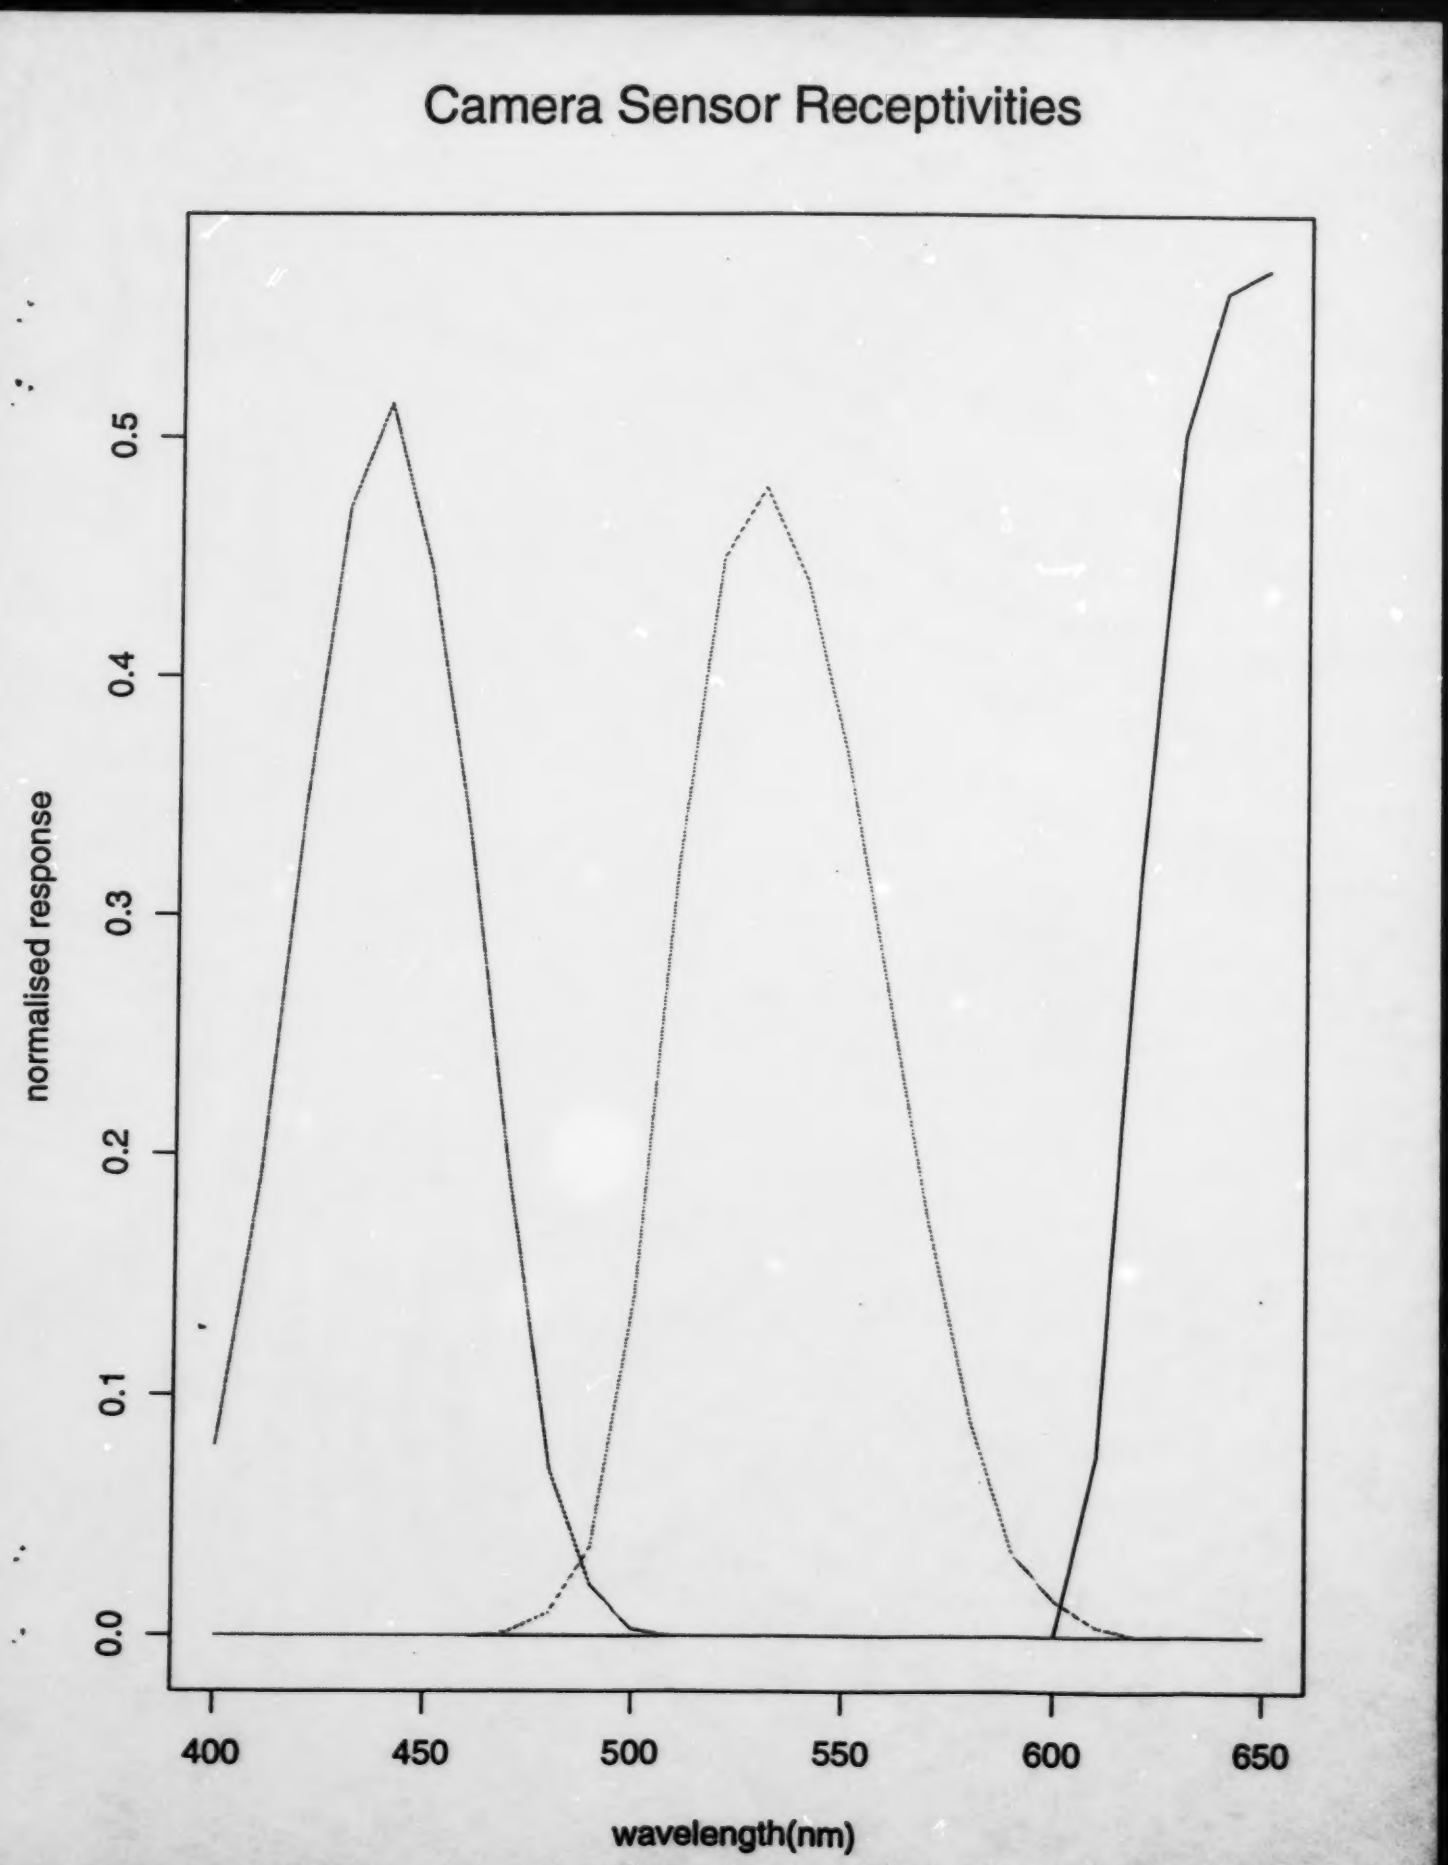
Camera Sensor Receptivities
normalised response
0.0
0.1
0.2
0.3
0.4
0.5
400
450
500
550
600
650
wavelength(nm)

Figure 2

Centre for Systems Science

• Simon Fraser University established the **Centre for Systems Science** as a provincial centre for advanced study. The objective of CSS is twofold: coordinate and combine advanced systems research at SFU, and encourage application of research results within industry in order to further the technological development of British Columbia.

• CSS has assembled three research groups to streamline its efforts. The orientation of these groups is as follows: intelligent systems, microelectronics, and computer and communications systems. Because of shared research interests, this latter group merges with LCCR (**Laboratory for Computer and Communications Research**) for the publication of technical reports.

• CSS's three groups involve over 80 skilled faculty researchers principally from SFU's Schools of Engineering Science and Computing Science, but also draw substantially on Mathematics, Physics, Geography, Philosophy, Linguistics, Psychology, Communication, and Kinesiology.